Where Does the Power Go and What to do About it?

James Hamilton

2008.12.07

Architect, Data Center Futures
e: JamesRH@microsoft.com
w: mvdirona.com/jrh/work
b: perspectives.mvdirona.com

Power & Related Costs Dominate

- **Assumptions:**
  - Facility: ~$200M for 15MW facility (15-year amort.)
  - Servers: ~$2k/each, roughly 50,000 (3-year amort.)
  - Average server power draw at 30% utilization: 75%
  - Commercial Power: ~$0.07/kWhr

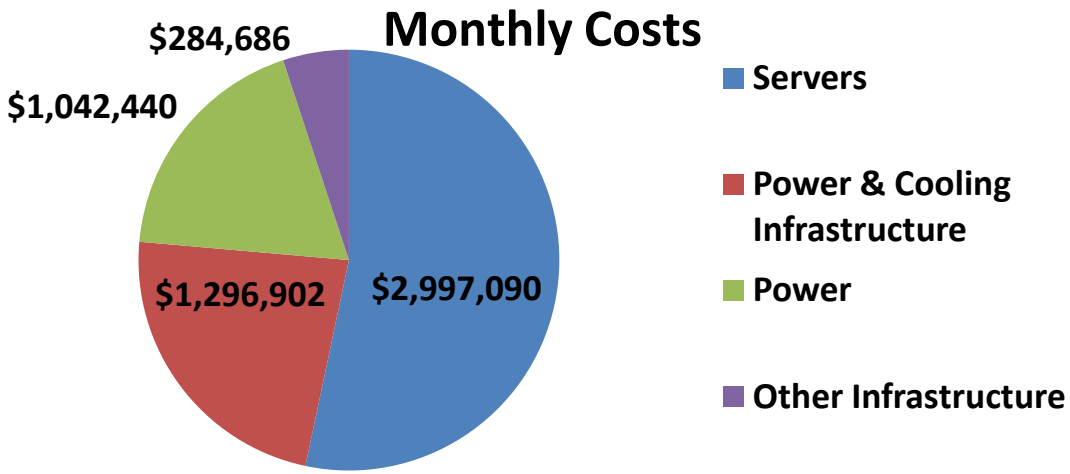

3yr server & 15 yr infrastructure amortization

- **Observations:**
  - $2.3M/month from charges functionally related to power
  - Power related costs trending flat or up while server costs trending down

Details at: http://perspectives.mvdirona.com/2008/11/28/CostOfPowerInLargeScaleDataCenters.aspx

Where Does the Power Go?

- **Assuming a pretty good data center with PUE ~1.7**
  - Each watt to server loses ~0.7W to power distribution losses & cooling
- **Power losses are easier to track than cooling:**
  - Power transmission & switching losses: 8%
    - Detailed power distribution losses on next slide
  - Cooling losses remainder:100-(59+8) => 33%
- **Data center power consumption:**
  - IT load (servers): 1/1.7=> 59%
  - Distribution Losses: 8%
  - Mechanical load(cooling): 33%

Power Distribution

115kv

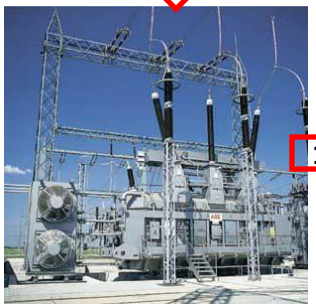

0.3% loss
99.7% efficient

13.2kv

8% distribution loss
.997^3*.94*.99 = 92.2%
4% to 5% loss attainable

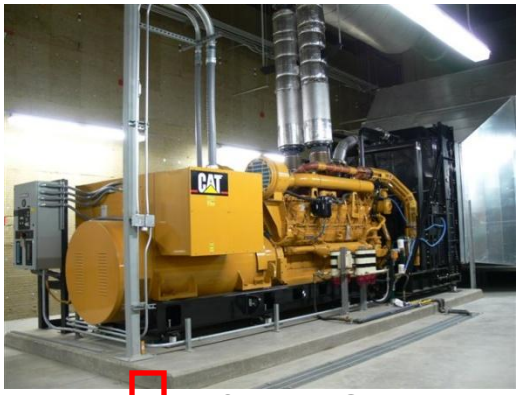

13.2kv

2.5MW Generator
~180 Gallons/hour

UPS:
Rotary or Battery

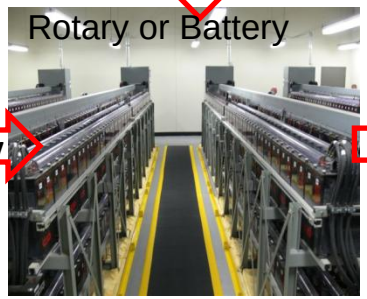

6% loss
94% efficient, >97% available

13.2kv

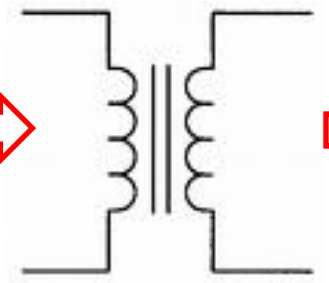

0.3% loss
99.7% efficient

480V

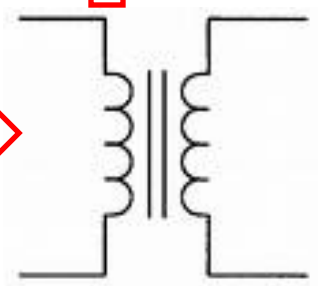

0.3% loss
99.7% efficient

208V

IT LOAD

~1% loss in switch Gear and conductors

Conventional Mechanical Design

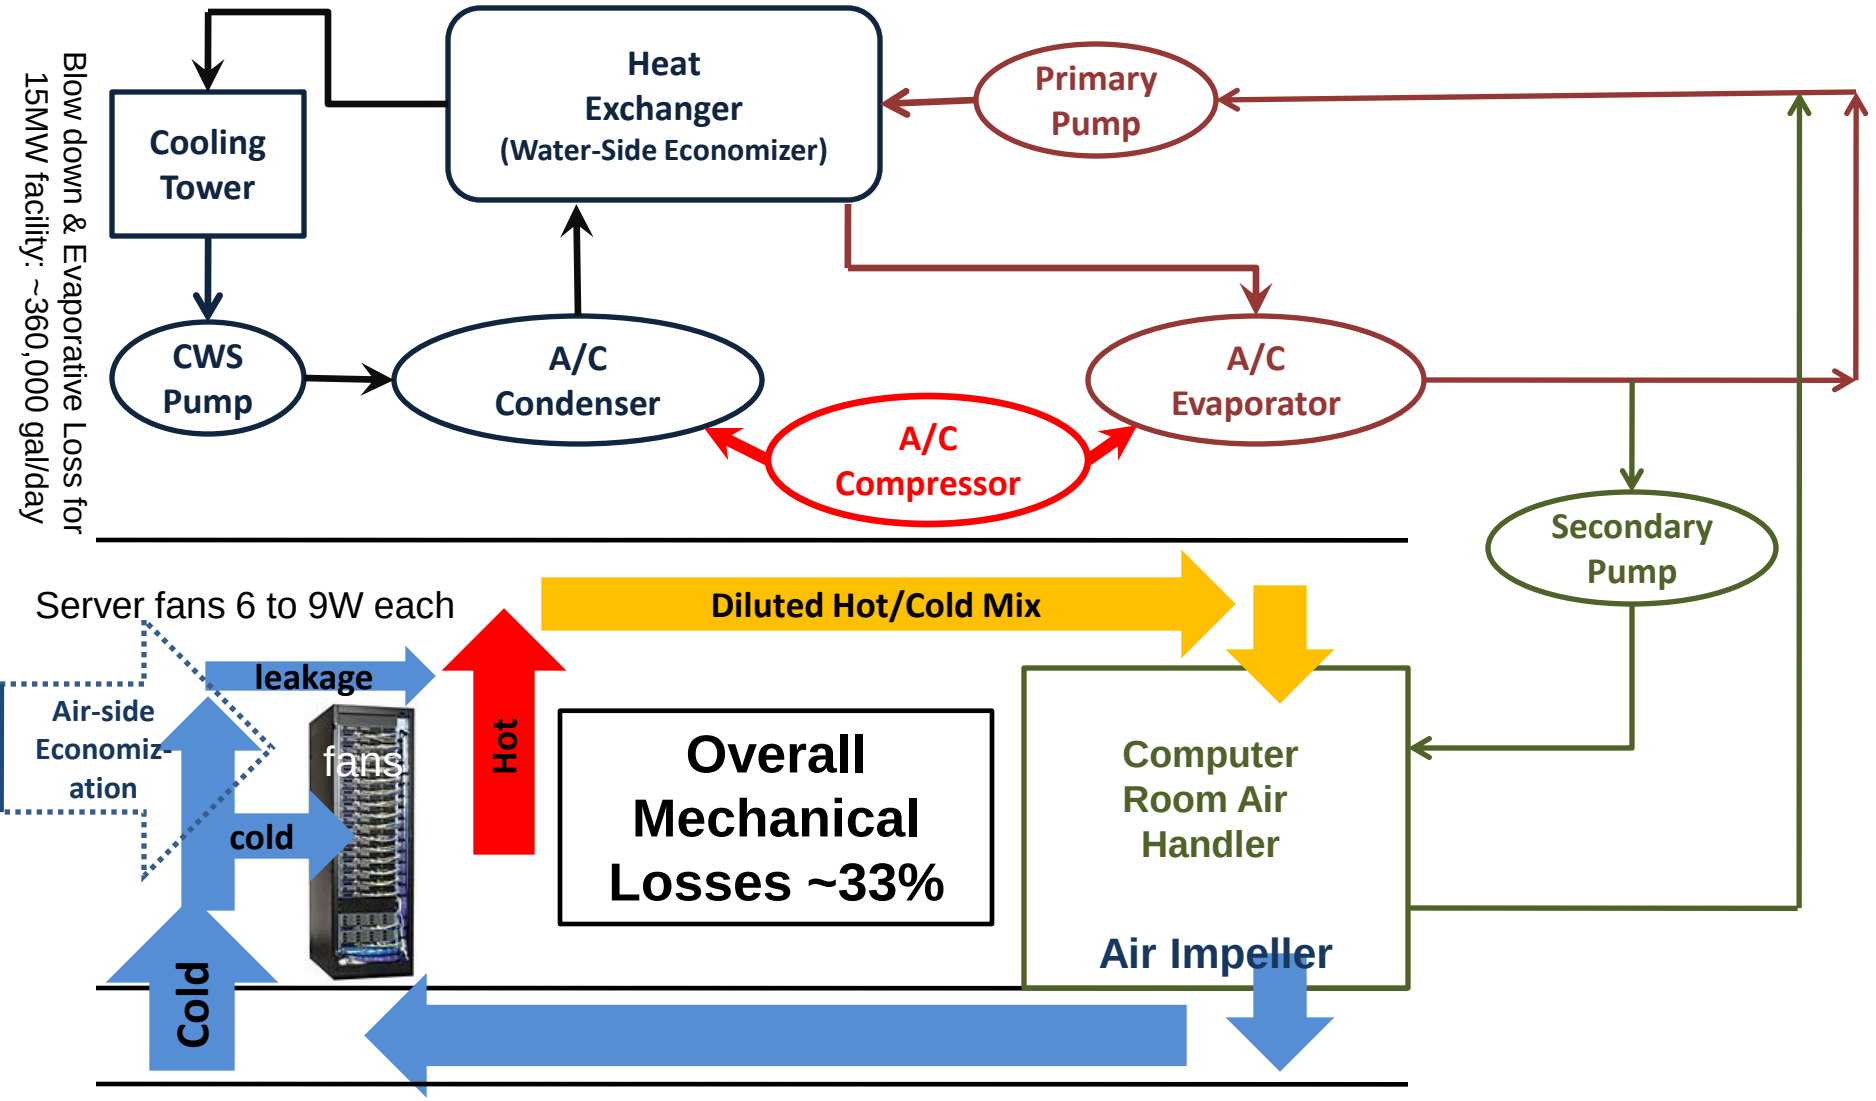

Cooperative Expendable Micro-Slice Servers

- CEMS: Cooperative Expendable Micro-Slice Servers
  - Correct system balance problem with less-capable CPU
    - Too many cores, running too fast, for memory, bus, disk, …
- Joint project with Rackable Systems (http://www.rackable.com/)

	System-X	CEMS V3 (Athlon 4850e)	CEMS V2 Athlon 3400e)	CEMS V1 (Athlon 2000+)
CPU load%	**56%**	**57%**	**57%**	**61%**
RPS	**95.92**	**75.26**	**54.27**	**17**
Price	**$2,371**	**$500**	**$685**	**$500**
Power	**295**	**60**	**39**	**33**
RPS/Price	0.04	0.15	0.08	0.03
RPS/Joule	0.32515254	1.254333333	1.391538462	0.515151515
RPS/Rack	1918.4	18062.4	13024.8	4080

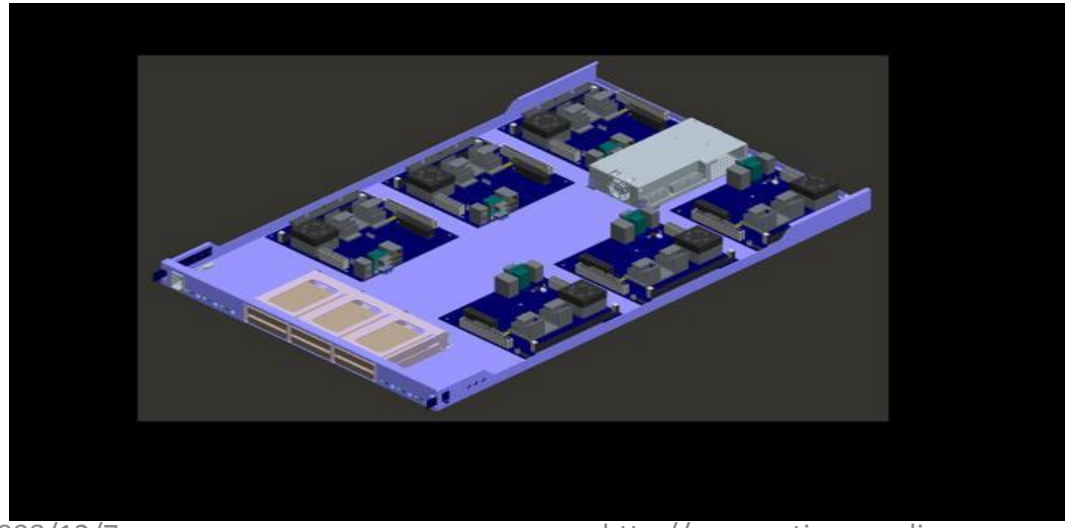

- CEMS V2 Comparison:
  - Work Done/$: +372%
  - Work Done/Joule +385%
  - Work Done/Rack: +941%

Update: New H/W SKU likely will improve numbers by factor of 2. CEMS still a win.

More Information

- **These slides**
  - Will be posted to http://mvdirona.com/jrh/work next week.
- **Designing & Deploying Internet-Scale Services**
  - http://mvdirona.com/jrh/talksAndPapers/JamesRH_Lisa.pdf
- **Architecture for Modular Data Centers**
  - http://mvdirona.com/jrh/talksAndPapers/JamesRH_CIDR.doc
- **Increasing DC Efficiency by 4x**
  - http://mvdirona.com/jrh/talksAndPapers/JamesRH_PowerSavings20080604.ppt
- **JamesRH Blog**
  - http://perspectives.mvdirona.com
- **Email**
  - JamesRH@microsoft.com or jrh@mvdirona.com